

BOUNDS ON THE DIAMETER OF CAYLEY GRAPHS OF THE SYMMETRIC GROUP

JOHN BAMBERG, NICK GILL, THOMAS P. HAYES, HARALD A. HELFGOTT, ÁKOS SERESS, AND PABLO SPIGA

ABSTRACT. In this paper we are concerned with the conjecture that, for any set of generators S of the symmetric group $\text{Sym}(n)$, the word length in terms of S of every permutation is bounded above by a polynomial of n . We prove this conjecture for sets of generators containing a permutation fixing at least 37% of the points.

1. INTRODUCTION

For a group G and a set S of generators of G , we write $\Gamma(G, S)$ for the *Cayley graph* of G with connection set S , that is, the graph with vertex set G and with edge set $\{\{g, sg\} \mid g \in G, s \in S\}$. The *diameter* $\text{diam}(\Gamma)$ of a graph Γ is the maximum distance among the vertices of Γ and, in the case of a Cayley graph $\Gamma(G, S)$, it is the maximum (over the group elements $g \in G$) of the shortest expression $g = s_1^{i_1} \cdots s_m^{i_m}$, with $s_k \in S$ and $i_k \in \{-1, 1\}$. We define the *diameter of a group G* as

$$\text{diam}(G) := \max\{ \text{diam}(\Gamma(G, S)) \mid S \text{ generates } G \}.$$

A first investigation of the diameter of Cayley graphs for general groups was undertaken by Erdős and Rényi [7]. Later Babai and Seress [4] obtained asymptotic estimates on $\text{diam}(G)$ depending heavily on the group structure of G . In particular the results in [4] highlight the discrepancy between the diameter of Cayley graphs of groups close to being abelian and the diameter of Cayley graphs of non-abelian simple groups. Moreover, [4] contains the following conjecture of Babai.

Conjecture 1.1 ([4, Conjecture 1.7]). *There exists $c > 0$ such that, for all non-abelian simple groups G , $\text{diam}(G) \leq (\log |G|)^c$.*

The conjecture remains open, although significant progress has been made. In particular, starting with the work of Helfgott on the groups $\text{PSL}(2, p)$ and $\text{PSL}(3, p)$ [11, 12] and based thereon, there has been a series of results [6, 8, 20] proving the conjecture for finite simple groups of Lie type of bounded rank. The best statement known at the time of writing is by Pyber and Szabó [20] and says that there exists a polynomial c such that, for a finite simple group G of Lie type of Lie rank r , we have $\text{diam}(G) \leq (\log |G|)^{c(r)}$. For the sake of comparison, Conjecture 1.1 asserts that c should be a constant rather than a polynomial.

The proofs of these theorems make use of new results in additive combinatorics, specifically on growth in simple groups. We note that the difficulties in generalizing these results from groups of bounded rank to those of unbounded rank seem closely related to difficulties in proving Conjecture 1.1 for the alternating groups $\text{Alt}(n)$. In both cases (that is, classical groups of unbounded rank and alternating groups) there are known counterexamples to general “growth results” for sets (see for example [19, 20, 23]), which were central to the approach used to prove Conjecture 1.1 for groups of Lie type of bounded rank. What is more, these two classes of counterexample are, in some sense, related.

In this paper we focus on the case where $G = \text{Alt}(n)$ or $\text{Sym}(n)$. Let Ω be a set of size n . For $g \in \text{Sym}(\Omega)$, define the *support* of g by $\text{supp}(g) = \{\gamma \in \Omega \mid \gamma^g \neq \gamma\}$. Observe that $\text{supp}(g)$ is equal to the complement in Ω of the *fixed set*, $\text{fix}(g)$, of g . Babai, Beals and Seress [2] proved the following result.

Theorem 1.2 ([2]). *For every $\varepsilon < 1/3$ there exists $c_\varepsilon > 0$ such that, if $G = \text{Sym}(n)$ or $\text{Alt}(n)$ and S is a set of generators of G containing an element g with $|\text{supp}(g)| \leq \varepsilon n$, then*

$$\text{diam}(\Gamma(G, S)) \leq c_\varepsilon n^8.$$

In this paper we provide a variant of the argument in [2] to prove the following stronger theorem.

2000 *Mathematics Subject Classification.* 20B25.

Key words and phrases. Cayley graph; diameter; Babai’s conjecture; Babai-Seress conjecture.

Theorem 1.3. *Let $C = 0.63$. There exists $c > 0$ such that, if $G = \text{Sym}(n)$ or $\text{Alt}(n)$ and S is a set of generators of G containing an element g with $|\text{supp}(g)| \leq Cn$, then*

$$\text{diam}(\Gamma(G, S)) \leq O(n^c).$$

We do not try to minimize the exponent c in the theorem. Our arguments give $c \leq 78$, but with some more work the bound on $\text{diam}(\Gamma(G, S))$ can be improved to at least $O(n^{66})$.

Theorem 1.3 also extends to directed graphs. Given $G = \langle S \rangle$, the *directed Cayley graph* $\vec{\Gamma}(G, S)$ is the graph with vertex set G and edge set $\{(g, sg) : g \in G, s \in S\}$. Analogously to the undirected case, the diameter of $\vec{\Gamma}(G, S)$ is defined as the maximum (taken over $g \in G$) of the shortest expression $g = s_1 \cdots s_m$, with each $s_k \in S$. By a theorem of Babai [1, Corollary 2.3], $\text{diam}(\vec{\Gamma}(G, S)) = O(\text{diam}(\Gamma(G, S)) \cdot (\log |G|)^2)$ holds for all groups G and sets S of generators, so we immediately obtain the following corollary.

Corollary 1.4. *Let $C = 0.63$. There exists $d > 0$ such that, if $G = \text{Sym}(n)$ or $\text{Alt}(n)$ and S is a set of generators of G containing an element g with $|\text{supp}(g)| \leq Cn$, then*

$$\text{diam}(\vec{\Gamma}(G, S)) \leq O(n^d).$$

We note that for arbitrary sets of generators the best known bound is quasipolynomial, by a recent result of Helfgott and Seress:

Theorem 1.5 ([13]). *For $G = \text{Alt}(n)$ and $\text{Sym}(n)$, $\text{diam}(G) = \exp(O((\log n)^4 \log \log n))$.*

The machinery developed in this paper turns out to have application to other questions within permutation group theory. Indeed it is possible to use variants of the results given in Section 3 to recover, and strengthen, classical results concerning multiply transitive groups due to Manning [16, 17, 18] and Wielandt [25]. This will be the subject of a forthcoming paper [10].

1.1. The main ideas. It is well-known and easy to see that if a set A of generators of $G = \text{Alt}(n)$ or $\text{Sym}(n)$ contains a 3-cycle t then every element of G can be written as a word of length less than n^4 in A . Indeed, repeatedly conjugating t by A gives all 3-cycles as words of length less than n^3 , and each element of $\text{Alt}(n)$ is a product of at most $\lfloor n/2 \rfloor$ 3-cycles. Finally, if $G = \text{Sym}(n)$ then one more multiplication by A gives words for all elements of G . Thus, given any set S of generators of G , in order to prove Conjecture 1.1, it is enough to construct a 3-cycle as a word in S of polynomial length.

We may try to reach a 3-cycle in stages, by constructing elements of smaller and smaller support. Up to very recently, the only subexponential method to obtain an element of support less than cn , for some constant $c < 1$, from arbitrary generating sets was in [3]. In that paper, iteration of the support reduction was utilized to prove $\text{diam}(G) = \exp((1 + o(1))\sqrt{n \log n})$, the only subexponential bound until Theorem 1.5.

Theorem 1.2 may be interpreted as a reduction of Conjecture 1.1 for alternating groups, to the problem of constructing an element g of moderately small support as a short word in an arbitrary set S of generators. The proof of Theorem 1.2 in [2] is based on the following observations.

- (BBS1) If $|\text{supp}(a)| < \varepsilon n$ for some $\varepsilon < 1/3$, $a \in G$, $G = \text{Alt}(n)$ or $\text{Sym}(n)$, and r is a random element of G then, for $b = ar$, the commutator $[a, b] = a^{-1}b^{-1}ab$ has support smaller than a with positive probability.
- (BBS2) In (BBS1), it is not necessary that r is a uniformly distributed random element of G . It is enough that, for some constant ℓ , r maps a sequence of distinct elements of length ℓ from the permutation domain nearly uniformly to all other sequences. Furthermore, random words r on any set S of generators of G , of length $n^{O(\ell)}$, satisfy this property.

In [2], the number $\ell = 3$ was chosen and a 3-cycle was constructed in $O(\log \log n)$ applications of (BBS1). A major conceptual novelty of [2] is that besides the natural action of G on n points and the action of G on itself as in $\Gamma(G, S)$, it is beneficial to work with other actions of G . This principle is more clearly formulated in [13]. In [2] and in the present paper, the action of G on sequences of length ℓ from the natural permutation domain is used, while in [13] other actions are utilized as well.

Chronologically, we made three improvements to the argument in [2].

- (NEW1) The conclusion of (BBS1) holds for $\varepsilon < 1/2$, implying a version of Theorem 1.3 with $C < 0.5$.
- (NEW2) With positive probability, the commutator $[a, b]$ has many fixed points, and also contains a significant number of 3-cycles. Thus, if $|\text{supp}(a)| < \varepsilon n$ for some $\varepsilon < 0.585$, then $[a, b]^3$ has support smaller than a . This implies Theorem 1.3 with $C = 0.585$.

(NEW3) With positive probability, the permutation $[a, b^{-1}][a, b]$ has many fixed points and 2-, 3-, 4- and 5-cycles. So, for $\varepsilon \leq 0.63$, $([a, b^{-1}][a, b])^{60}$ has support smaller than a .

In this paper we only prove the strongest version based on (NEW3). We have to overcome several technical difficulties: (i) the analysis of the local behaviour (i.e., finding how a and b should interact on small subsets Δ of the natural domain such that $[a, b^{-1}][a, b]$ forms a short cycle on some points of Δ); (ii) ensuring that $([a, b^{-1}][a, b])^{60}$ is not the identity of G ; and (iii) handling the special case when the originally given generator a has order $2^x 3^y$ for some $x, y \geq 0$. We shall apply the argument of (BBS2) with $\ell = 26$.

The structure of this paper is as follows. In Section 2, we collect basic concepts regarding groups and graphs, and introduce the central notion of $\alpha\beta$ -trees. These are the objects describing the possible local interactions of a and b . We also introduce our probabilistic method. In Section 3 we give a graph-theoretic technique for estimating the number of fixed points of $w(a, b)$, where $w(\alpha, \beta)$ is a reduced word in α and β , and a and b are particular conjugate permutations of $\text{Sym}(\Omega)$. In Section 4, we apply the results of Section 3 to the word $[\alpha, \beta^{-1}][\alpha, \beta] = \alpha^{-1}\beta\alpha\beta^{-1}\alpha^{-1}\beta^{-1}\alpha\beta$ and we prove Theorem 1.3. In Section 5 we discuss some possible extensions of Theorem 1.3.

1.2. Acknowledgements. All authors would like to thank Gordon Royle for useful discussions in the initial phase of this research. In addition expenses for N.G. and H.H. to visit the University of Western Australia were paid for by a UWA Research Collaboration Award which was awarded to a team including Prof. Royle. We are, therefore, doubly grateful to Prof. Royle for without this financial support it is unlikely that this research would have been undertaken.

A.S is supported in part by the NSF and by ARC Grant DP1096525. N.G. was a frequent visitor to the University of Bristol during the period of this research and is grateful for the generous support he received from the maths department there.

2. BASIC CONCEPTS

In this section we collect definitions and basic results that will be needed in the proof of Theorem 1.3.

2.1. Permutation groups. Let $\Omega = \{1, 2, \dots, n\}$. We use $\text{Sym}(n)$ and $\text{Sym}(\Omega)$ interchangeably; more exactly, we use $\text{Sym}(\Omega)$ when we emphasise the action of $\text{Sym}(n)$ on Ω . Let $S \subseteq \text{Sym}(n)$ and $k, l \in \mathbb{Z}^+$. Define

$$S^\ell = \{s_1 \cdots s_\ell \mid s_1, \dots, s_\ell \in S\}; \quad S^{-1} = \{s^{-1} \mid s \in S\}.$$

For $\omega \in \Omega$ and $a, g \in \text{Sym}(n)$, we write ω^g for the image of ω under g ; and a^g for $g^{-1}ag$. We denote by $\Omega_{(k)}$ the set of k -tuples of distinct elements of Ω and we write $n_{(k)} = |\Omega_{(k)}| = n(n-1) \cdots (n-k+1)$.

We shall use the following result of J. Whiston [24].

Lemma 2.1 ([24]). *Any set S of generators for $G = \text{Sym}(n)$ or $\text{Alt}(n)$ contains a subset A of cardinality less than or equal to $n-1$ that also generates G .*

If A and B are two sets of generators for G with $A \subseteq B$ then clearly $\text{diam}(\Gamma(G, B)) \leq \text{diam}(\Gamma(G, A))$. Therefore, Lemma 2.1 implies that it is enough to prove Theorem 1.3 for sets S of generators of size at most n which contain an element g of small support.

2.2. Graphs. In this paper a *graph* X is a finite connected *directed* graph. Moreover, we allow loops on the vertices of X and multiple edges. We do not assume that X is *strongly connected*, i.e., it is possible that there is no directed path between some vertices x and y of X . We write $V(X)$ for the set of vertices of X and $E(X)$ for the set of edges of X . An edge e running from vertex i to vertex j will be written (i, j) but we warn that this notation is ambiguous as there may be more than one such edge.

We define an $\alpha\beta$ -graph, T say, to be a graph together with a label, α or β , attached to every edge. We require that for each vertex $v \in V(T)$ and for each $\gamma \in \{\alpha, \beta\}$, v is incident with at least one edge labelled by γ .¹ We also require that at most one edge starting at v is labelled by γ and at most one edge ending at v is labelled by γ . Here a loop at v counts as one incoming and one outgoing edge for v . Notice that all vertices of an $\alpha\beta$ -graph have in-degree at most 2 and out-degree at most 2. For $\gamma \in \{\alpha, \beta\}$, we define T_γ to be the subgraph of T with vertex set $V(T)$ and edge set the set of edges of T labelled by γ .

¹We impose this condition because it is necessary in some of the arguments used in Section 3. It is not *a priori* necessary to our strategy for analysing the fixed points of words; this condition limits our implementation to those reduced words w on $\{\alpha, \beta, \alpha^{-1}, \beta^{-1}\}$ for which all exponents are equal to ± 1 .

We say that a cycle C of T is *monochromatic* if all of its edges are labelled α (resp. β), that is, C is a sequence of vertices $(v_1, \dots, v_r, v_{r+1})$ with $v_{r+1} = v_1$ and $r \geq 2$, and where for each $i \in \{1, \dots, r\}$ the ordered pair (v_i, v_{i+1}) is an edge of T labelled α (resp. β).

Given permutations $a, b \in \text{Sym}(\Omega)$ and an injective map $\iota : V(T) \rightarrow \Omega$, we say that T_α is *hosted by* (ι, a) if $(x\iota)^a = y\iota$ for each edge or loop $(x, y) \in E(T_\alpha)$. Similarly, T_β is *hosted by* (ι, b) if $(x\iota)^b = y\iota$ for each edge or loop $(x, y) \in E(T_\beta)$. Finally, T is *hosted by* (ι, a, b) if T_α is hosted by (ι, a) and T_β is hosted by (ι, b) .

Let T be an $\alpha\beta$ -graph. Observe that for $\gamma \in \{\alpha, \beta\}$, the connected components of T_γ are of three types:

- (1) γ -*loops*: isolated vertices, namely the vertices v of T having a loop at v labelled with γ ;
- (2) γ -*cycles*: monochromatic cycles all of whose edges are labelled γ ;
- (3) γ -*paths*: maximal directed paths such that all edges are labelled with γ .

We denote by $l_\gamma(T)$ the number of γ -loops and by $p_\gamma(T)$ the number of γ -paths and γ -cycles.

We say that an $\alpha\beta$ -graph is an $\alpha\beta$ -*tree* if all undirected cycles are monochromatic (and so necessarily they are also directed cycles). Note that an $\alpha\beta$ -tree may not be a tree in the usual graph-theoretic sense. The following result will be crucial.

Lemma 2.2. *Let T be an $\alpha\beta$ -graph. Then $p_\alpha(T) + p_\beta(T) + l_\alpha(T) + l_\beta(T) \leq |V(T)| + 1$. Moreover, equality holds if and only if T is an $\alpha\beta$ -tree.*

Proof. Let B be the graph with vertex set the set of α -loops, α -paths, α -cycles, β -loops, β -paths and β -cycles of T . We declare two distinct vertices x and y of B adjacent if there exists a vertex v of T such that v is incident with both x and y . By construction, B is bipartite (with the α -objects comprising one class of the bipartition and the β -objects the other) and $|V(B)| = p_\alpha(T) + p_\beta(T) + l_\alpha(T) + l_\beta(T)$. By the definition of $\alpha\beta$ -graphs, each vertex v of T is incident with exactly one component of T_α and with exactly one component of T_β . Hence v defines exactly one edge in B and so $|E(B)| = |V(T)|$. Since T is connected, the graph B is connected and so $|V(B)| \leq |E(B)| + 1 = |V(T)| + 1$, proving the first claim.





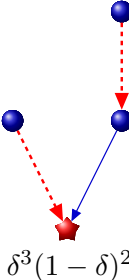
Observe that if T contains a non-monochromatic cycle, then B contains a cycle. Conversely if B contains a cycle, then T must contain a non-monochromatic cycle. We conclude that B is a tree if and only if T has no non-monochromatic cycles, and the second claim follows from the standard fact that B is a tree if and only if $|V(B)| = |E(B)| + 1$. \square

We end this subsection with some more definitions. For an $\alpha\beta$ -graph T and for $0 < \delta < 1$, we define $\delta_T := (1 - \delta)^{l_\alpha(T) + l_\beta(T)} \delta^{p_\alpha(T) + p_\beta(T)}$. An *isomorphism* between $\alpha\beta$ -graphs T_1 and T_2 is defined to be a bijection $\varphi : V(T_1) \rightarrow V(T_2)$ that preserves edges and edge-labels. If, for some $x, y \in V(T_1)$, there are two directed edges from x to y in T_1 (with necessarily different labels) then in T_2 there are also two directed edges from $x\varphi$ to $y\varphi$. We define an *automorphism* of T to be an isomorphism between T and itself. The set of all such permutations is the *automorphism group* $\text{Aut}(T)$.

2.3. Words and graphs. Let T be an $\alpha\beta$ -graph and $w = w_1 w_2 \dots w_k$ be a reduced word with $w_1, \dots, w_k \in \{\alpha, \alpha^{-1}, \beta, \beta^{-1}\}$. We say that T *admits* w , if there exists a vertex x of T such that by starting at x and by tracing the edges and the loops of T with labels (w_1, w_2, \dots, w_k) , we visit *all* vertices and edges of T and we return to the vertex x (here by abuse of notation, we interpret the label α^{-1} as the label α with the edge pointing in the opposite direction, and a similar convention holds for β). The vertex $x \in V(T)$ is called a *fixed vertex* for (T, w) ; note that there may be more than one such vertex in T .

The following Table 1 contains pairwise non-isomorphic $\alpha\beta$ -trees that admit the word $w = [\alpha, \beta^{-1}][\alpha, \beta]$. Here dashed (red) lines are labelled with α , solid (blue) lines are labelled with β , and for simplicity of drawing all loops are omitted; thus, any vertex which is not incident to a red line (or blue line, respectively) is in fact incident to an α -loop (or β -loop, respectively). The fixed vertices are written as red stars, and under each graph T we have written the value for δ_T .

TABLE 1. $\alpha\beta$ -trees admitting $w = [\alpha, \beta^{-1}][\alpha, \beta] = \alpha^{-1}\beta\alpha\beta^{-1}\alpha^{-1}\beta^{-1}\alpha\beta$.

				
$(1 - \delta)^2$	$\delta^1(1 - \delta)^2$	$\delta^1(1 - \delta)^3$	$\delta^3(1 - \delta)^2$	$\delta^3(1 - \delta)^2$

Next, we explain how $\alpha\beta$ -trees can be used to estimate the number of fixed points in certain permutations. Let T be an $\alpha\beta$ -tree admitting the reduced word $w = w_1w_2 \cdots w_k$, and let x be a fixed vertex of (T, w) . Starting at x and tracing w we obtain a sequence $U = (x = x_0, x_1, \dots, x_k = x)$ of vertices of T such that, by definition, all $v \in V(T)$ occur in U .

Let $a, b \in \text{Sym}(\Omega)$. If T is hosted by (ι, a, b) for some injective map $\iota : V(T) \rightarrow \Omega$ then, starting at $x\iota$ and tracing the word w by using a and b for the labels α and β , respectively, we obtain the sequence $U\iota = (x_0\iota, x_1\iota, \dots, x_k\iota)$. In particular, $x\iota$ and w uniquely determine the entire map ι and $x\iota$ is a fixed point of the permutation $w(a, b)$.

For fixed a, b , and w as above, let T_1, \dots, T_m be pairwise non-isomorphic $\alpha\beta$ -trees admitting w . For $1 \leq j \leq m$, we denote the number of fixed vertices in T_j (with respect to w) by $\text{fixed}(T_j)$. Also, let I_j be an index set such that for $z \in I_j$, there exists $\iota_z : V(T_j) \rightarrow \Omega$ with T_j hosted by (ι_z, a, b) .

Lemma 2.3. *Let a, b, w, m, T_j, I_j be as in the previous paragraph.*

- (i) *If x is a fixed vertex of (T_{j_1}, w) , y is a fixed vertex of (T_{j_2}, w) , and $x\iota_{z_1} = y\iota_{z_2}$ for some $z_1 \in I_{j_1}$ and $z_2 \in I_{j_2}$ then $j_1 = j_2$ and x, y are in the same orbit of $\text{Aut}(T_{j_1})$.*
- (ii) *The number of fixed points of $w(a, b)$ is at least*

$$|\text{fix}(w(a, b))| \geq \sum_{j=1}^m \frac{|I_j| \cdot \text{fixed}(T_j)}{|\text{Aut}(T_j)|}.$$

Proof. (i) Since $x\iota_{z_1} = y\iota_{z_2}$ and w determines $T_{j_1}\iota_{z_1} = T_{j_2}\iota_{z_2}$, the map $\iota_{z_1}\iota_{z_2}^{-1}$ is a label-preserving isomorphism between T_{j_1} and T_{j_2} . Therefore, $j_1 = j_2$ and $\iota_{z_1}\iota_{z_2}^{-1} \in \text{Aut}(T_{j_1})$. Moreover, since $x\iota_{z_1}\iota_{z_2}^{-1} = y$, x and y are in the same orbit of $\text{Aut}(T_{j_1})$.

(ii) For $1 \leq j \leq m$, let F_j be the set of fixed points of $w(a, b)$ of the form $x\iota_z$, for some fixed vertex $x \in V(T_j)$ and $z \in I_j$. By (i), the sets F_j are pairwise disjoint. Moreover, for $z \in I_j$, ι_z contributes $\text{fixed}(T_j)$ fixed points to F_j , giving a total count of $|I_j| \cdot \text{fixed}(T_j)$. Part (i) also implies that any element of F_j occurs in this count at most $|\text{Aut}(T_j)|$ times. \square

2.4. Walks on graphs. In this subsection we consider graphs X that are symmetric and regular in the following sense. *Symmetric* means that for any two vertices $x, y \in V(X)$, the number of edges from x to y is the same as the number of edges from y to x . *Regular* of valency d means that each $x \in V(X)$ has in-degree and out-degree d . For $x \in V(X)$, we denote by $\Delta(x)$ the d -element multiset $\{y \mid (x, y) \in E(X)\}$.

Definition 2.4. A *lazy random walk* on X is a discrete stochastic process where a particle moves from vertex to vertex in X . If, after k steps, the particle is at $x \in V(X)$ and $\Delta(x) = \{y_1, \dots, y_d\}$ then at step $k + 1$ the particle

- (i) stays at x with probability $1/2$;
- (ii) moves to vertex y_i with probability $1/(2d)$, for all $i = 1, \dots, d$.

The asymptotic rate of convergence for the probability distribution of a particle in a lazy random walk on X is an important and well-studied problem in combinatorics and computer science (see [15]). For $x, y \in V(X)$, we write $p_k(x, y)$ for the probability that the particle is at vertex y after k steps of a lazy random walk starting at x . For a fixed $\varepsilon > 0$, the *mixing time for ε* is the minimum value of k such that

$$\frac{1}{|V(X)|}(1 - \varepsilon) \leq p_k(x, y) \leq \frac{1}{|V(X)|}(1 + \varepsilon)$$

for all $x, y \in V(X)$. The following estimate is well-known; for a proof, see e.g. [13, Section 4].

Lemma 2.5. *Let X be a connected, symmetric, and regular directed graph of valency d and with N vertices, and let $\varepsilon > 0$. Then the mixing time for ε is at most $N^2 d \log(N/\varepsilon)$.*

For $G = \text{Sym}(\Omega)$ or $\text{Alt}(\Omega)$ and $G = \langle S \rangle$, we are interested in the following symmetric and regular directed graphs X_k , for positive integers k . Let $V(X_k) := \Omega_{(k)}$, and $E(X_k) := \{(x, x^g) \mid x \in \Omega_{(k)} \text{ and } g \in S \cup S^{-1}\}$. Clearly X_k has $n_{(k)}$ vertices, is connected, is symmetric, and is regular of valency $|S \cup S^{-1}|$.

It is useful to induce random walks on the graphs X_k for different k at the same time. This is done as follows. First, we choose a subset $J \subseteq \{1, 2, \dots, \ell\}$, where J is the set of steps when the particle moves to a neighbour of the current position as in Definition 2.4(ii). The length $j := |J|$ is chosen from the binomial distribution $B(\ell, 1/2)$ and then J itself is chosen from the uniform distribution on the j -element subsets of $\{1, 2, \dots, \ell\}$. Finally, for $i \in J$, we choose $g_i \in S \cup S^{-1}$ uniformly and use g_i in the i^{th} step to define the edge on which the particle moves. The overall effect, that is, the trajectory of a lazy random walk with initial position $x \in V(X_k)$, is the same as computing the image of x under the permutation $r = \prod_{i \in J} g_i$; we say that the permutation r is realised by the lazy random walk. The construction of r uses only the number ℓ and $S \cup S^{-1}$, so the permutation r can be considered as realised by lazy random walks in more than one graph X_k . Of course, these lazy random walks are not independent.

Lemma 2.5 will be useful to us in the following form.

Lemma 2.6. *Let S be a set of generators of $\text{Sym}(\Omega)$ or $\text{Alt}(\Omega)$ of cardinality at most n , and let k be a positive integer. Fix $0 < \varepsilon < 1$ and set $\ell \geq 2n^{2k+1} \log(n^k/\varepsilon)$. Then, if $r \in \text{Sym}(\Omega)$ is realised by a lazy random walk of length ℓ on X_k , and $x, y \in \Omega_{(k)}$, then*

$$(1 - \varepsilon) \frac{1}{n_{(k)}} \leq \mathbb{P}(x^r = y) \leq (1 + \varepsilon) \frac{1}{n_{(k)}}.$$

Proof. Recall that $n_{(k)} = |V(X_k)|$ by definition, and $n_{(k)} \leq n^k$. Furthermore $|S \cup S^{-1}| \leq 2n$ and the proof follows from Lemma 2.5. \square

Note that if $r \in \text{Sym}(\Omega)$ is realised by a lazy random walk of length ℓ on X_k , then $r \in (S \cup S^{-1} \cup \{1\})^\ell$.

3. PRIMARY MACHINERY

The results of this section provide the primary machinery for a proof of Theorem 1.3. The main step of the proof is that given a set S of generators for $G = \text{Sym}(n)$ or $\text{Alt}(n)$ and $a \in S$ of support size $|\text{supp}(a)| = \delta n$, we would like to construct a permutation as a short word in S with support size less than δn . The permutations we consider are of the form $w(a, a^r)$, for an appropriately chosen reduced word w in the symbols $\{\alpha, \alpha^{-1}, \beta, \beta^{-1}\}$. In this section, we assume that w is given, and describe how to choose $r \in G$ such that $w(a, a^r)$ has many fixed points. We obtain r as a permutation realised by a lazy random walk.

Let $w = w_1 w_2 \dots w_k$ and let $\mathcal{T} = \{T_1, \dots, T_m\}$ be a set of pairwise non-isomorphic $\alpha\beta$ -trees admitting w . By Lemma 2.3(ii), we would like to choose r so that each $T \in \mathcal{T}$ is hosted by (ι, a, a^r) for many maps $\iota : V(T) \rightarrow \Omega$. As a is fixed, it is beneficial first to examine embeddings of T_α and T_β separately.

We prove results for two kinds of permutations a . In the “generic” case, all non-trivial cycles of a are long, compared to the α - and β -paths occurring in trees T and in the “special” case $\text{supp}(a)$ consists of short cycles of equal length. We fix a small set $\Lambda \subset \Omega$ (in the application in Section 4, $|\Lambda| \leq 10$) and require that r fixes Λ setwise and acts on Λ on some prescribed way. The purpose of prescribing the action of r on a small set is to ensure that $w(a, a^r)$ is not trivial. (This trick has already been used in [2].) As the points in Λ play a special role, we are only interested in injections $\iota : V(T) \rightarrow \Omega \setminus \Lambda$. First, we handle the “generic” case.

Lemma 3.1. *Let $0 < \delta_0 < 1/2$ and let κ, λ, N be positive integers. Suppose that $a \in \text{Sym}(\Omega)$ has no cycles of length less than N and that $|\text{supp}(a)| = \delta n$, for some $\delta \in (\delta_0, 1 - \delta_0)$. Let $\gamma \in \{\alpha, \beta\}$ and let T be an $\alpha\beta$ -tree such that $|V(T)| \leq \kappa$, T has no γ -cycles, and every γ -path in T has at most N vertices. Let $\Lambda \subseteq \Omega$, $|\Lambda| \leq \lambda$, and let*

$$\mathcal{S}_\gamma(T) := \{\iota : V(T) \rightarrow \Omega \setminus \Lambda \mid T_\gamma \text{ is hosted by } (\iota, a)\}.$$

Then

$$|\mathcal{S}_\gamma(T)| \geq C(\delta_0, \kappa, \lambda, N, n) (1 - \delta)^{l_\gamma(T)} \delta^{p_\gamma(T)} n^{l_\gamma(T) + p_\gamma(T)},$$

where $C(\delta_0, \kappa, \lambda, N, n)$ is a function with $\lim_{n \rightarrow \infty} C(\delta_0, \kappa, \lambda, N, n) = 1$.

By $\lim_{n \rightarrow \infty} C(\delta_0, \kappa, \lambda, N, n)$, we mean that the variables $\delta_0, \kappa, \lambda$ and N are fixed, and n goes to ∞ .

Proof. Let $v_1, \dots, v_{l_\gamma(T)}$ be the loops of T_γ and let $P_1, \dots, P_{p_\gamma(T)}$ be the directed paths of T_γ . We embed the components of T_γ into $\Omega \setminus \Lambda$ one-by-one, and estimate $|\mathcal{S}_\gamma(T)|$ by the product of the number of possible embeddings at each step.

The vertices v_i , for $1 \leq i \leq l_\gamma(T)$, have to be mapped to fixed points of a . As a has at least $n - \delta n - \lambda$ fixed points outside Λ , $v_1, \dots, v_{l_\gamma(T)}$ can be chosen in at least $(n - \delta n - \lambda)(n - \delta n - \lambda - 1) \cdots (n - \delta n - \lambda - l_\gamma(T) + 1) \geq (n - \delta n - \lambda - \kappa)^{l_\gamma(T)}$ distinct ways.

Next, we consider the directed paths in T_γ . For $1 \leq i \leq p_\gamma(T)$, we write $P_i = (w_{i,0}, \dots, w_{i,c_i})$. Now, once the image of $w_{i,0}$ under ι is chosen, in order to guarantee that ι hosts the path P_i , we require that $w_{i,j}\iota = (w_{i,0}\iota)^{a^j}$ for each $j = 0, \dots, c_i$. Hence, for each $i \in \{1, \dots, p_\gamma(T)\}$, the image of P_i under ι is uniquely determined by $w_{i,0}\iota$. Let Δ_i be the union of the sets $P_z\iota$, for $z < i$. Since by hypothesis a has no cycles of length $N - 1$ or shorter and since T has no γ -path of length greater than N , the only requirement for choosing the image of $w_{i,0}$ under ι is that

$$w_{i,0}\iota \notin \bigcup_{j=0}^{c_i} (\Delta_i \cup \Lambda)^{a^{-j}}$$

(since we have to avoid Λ and the ι -images of previously mapped P_z). A gross overestimate for the size of this union is $(c_i + 1)(\lambda + |V(T)|) \leq \kappa\lambda + \kappa^2$, and so $w_{i,0}\iota$ can be chosen in at least $\delta n - \kappa\lambda - \kappa^2$ ways. Summarizing, we obtain

$$|\mathcal{S}_\gamma(T)| \geq (n - \delta n - \lambda - \kappa)^{l_\gamma(T)} (\delta n - \kappa\lambda - \kappa^2)^{p_\gamma(T)}.$$

By factoring out δ , $(1 - \delta)$, and n , we get

$$|\mathcal{S}_\gamma(T)| \geq C(\delta_0, \kappa, \lambda, N, n) (1 - \delta)^{l_\gamma(T)} \delta^{p_\gamma(T)} n^{l_\gamma(T) + p_\gamma(T)},$$

where

$$C(\delta_0, \kappa, \lambda, N, n) = \left(1 - \frac{\kappa + \lambda}{\delta_0 n}\right)^\kappa \left(1 - \frac{\kappa\lambda + \kappa^2}{\delta_0 n}\right)^\kappa.$$

Clearly, $\lim_{n \rightarrow \infty} C(\delta_0, \kappa, \lambda, N, n) = 1$. □

The case of “special” permutations a is very similar.

Lemma 3.2. *Let $0 < \delta_0 < 1/2$ and let κ, λ, N be positive integers. Suppose that every cycle of $a \in \text{Sym}(\Omega)$ has length 1 or N and $|\text{supp}(a)| = \delta n$, for some $\delta \in (\delta_0, 1 - \delta_0)$. Let $\gamma \in \{\alpha, \beta\}$ and let T be an $\alpha\beta$ -tree such that $|V(T)| \leq \kappa$, every γ -cycle in T has N vertices, and every γ -path in T has at most N vertices. Let $\Lambda \subseteq \Omega$, $|\Lambda| \leq \lambda$, and let*

$$\mathcal{S}_\gamma(T) := \{\iota : V(T) \rightarrow \Omega \setminus \Lambda \mid T_\gamma \text{ is hosted by } (\iota, a)\}.$$

Then

$$|\mathcal{S}_\gamma(T)| \geq C(\delta_0, \kappa, \lambda, N, n) (1 - \delta)^{l_\gamma(T)} \delta^{p_\gamma(T)} n^{l_\gamma(T) + p_\gamma(T)},$$

where $C(\delta_0, \kappa, \lambda, N, n)$ is a function with $\lim_{n \rightarrow \infty} C(\delta_0, \kappa, \lambda, N, n) = 1$.

Proof. We may follow almost verbatim the proof of Lemma 3.1. The only difference is that the list $P_1, \dots, P_{p_\gamma(T)}$ may also contain γ -cycles, so we have to change slightly the definition of the vertices $w_{i,0}$. If P_i is a γ -path then $w_{i,0}$ is the starting vertex of the path as before, while if P_i is a γ -cycle then $w_{i,0}$ can be chosen as an arbitrary vertex of P_i . Since the γ -cycles of T_γ have the same length as the cycles in $\text{supp}(a)$, the rest of the proof goes through without any modification. □

Now we are ready to prove the main result of this section. Let $0 < \delta_0 < 1/2$ and $\kappa, \lambda, N \in \mathbb{Z}^+$ be fixed, and let $w = w_1 \cdots w_k$ be a reduced word in $\{\alpha, \alpha^{-1}, \beta, \beta^{-1}\}$. Suppose further that $\mathcal{T} = \{T_1, \dots, T_m\}$ is a set of $\alpha\beta$ -trees admitting w and $|V(T)| \leq \kappa$ for all $T \in \mathcal{T}$. Let $G = \text{Sym}(n)$ or $\text{Alt}(n)$ be generated by a set S of cardinality at most n . We do not have to distinguish the two cases (“generic” and “special”) for a anymore, so let $a \in \text{Sym}(\Omega)$ with $|\text{supp}(a)| = \delta n$ for some $\delta \in (\delta_0, 1 - \delta_0)$ and suppose that either

- all non-trivial cycles in a have length at least N , none of the $\alpha\beta$ -trees $T \in \mathcal{T}$ have any cycles, and every α - and β -path in T has at most N vertices; or
- every cycle of a has length 1 or N , for all $T \in \mathcal{T}$ every α - and β -cycle has N vertices, and every α - and β -path has at most N vertices.

Let $\Lambda \subseteq \Omega$, $|\Lambda| \leq \lambda$, let $g \in \text{Sym}(\Lambda)$, and let $S_\gamma(T)$ be as in Lemmas 3.1 and 3.2. Finally, let an error bound $\varepsilon > 0$ be given. We “collect” errors in estimates from different sources, so we choose $\varepsilon' < \varepsilon$ such that

$$\frac{(1 - \varepsilon')^3}{1 + \varepsilon'} = 1 - \varepsilon.$$

Moreover, we may assume that n is larger than a bound $n_0(\delta_0, \kappa, \lambda, N, \varepsilon)$ depending only on $\delta_0, \kappa, \lambda, N$, and ε such that:

(i) For all $\gamma \in \{\alpha, \beta\}$ and for all $T \in \mathcal{T}$,

$$(3.2.1) \quad |\mathcal{S}_\gamma(T)| > (1 - \varepsilon') (1 - \delta)^{l_\gamma(T)} \delta^{p_\gamma(T)} n^{l_\gamma(T) + p_\gamma(T)}.$$

(Note that this inequality is satisfied by large enough n , by Lemmas 3.1 and 3.2.)

(ii) $n^{2(\kappa + \lambda + 1)} > 2n^{2(\kappa + \lambda) + 1} \log(n^{\kappa + \lambda} / \varepsilon')$.

Recall that $\delta_T = (1 - \delta)^{l_\alpha(T) + l_\beta(T)} \delta^{p_\alpha(T) + p_\beta(T)}$ and that $\text{fixed}(T)$ denotes the number of fixed vertices of (T, w) .

Theorem 3.3. *With the notation of the previous paragraph, there exists $r \in \text{Sym}(\Omega)$ realised by a lazy random walk of length $n^{2(\kappa + \lambda + 1)}$ such that $r|_\Lambda = g$ and*

$$|\text{fix}(w(a, a^r))| > (1 - \varepsilon) n \sum_{j=1}^m \frac{\delta_{T_j} \cdot \text{fixed}(T_j)}{|\text{Aut}(T_j)|}.$$

By $r|_\Lambda$ we mean the restriction of the permutation r (considered as a function $r : \Omega \rightarrow \Omega$) to Λ .

Proof. Let r be realised by a lazy random walk of length $\ell := n^{2(\kappa + \lambda + 1)}$. Our main goal is to give an estimate for the conditional expectation $\mathbb{E}(|\text{fix}(w(a, a^r))| \mid r|_\Lambda = g)$.

Let $T \in \mathcal{T}$ and $\iota \in \mathcal{S}_\alpha(T)$ be arbitrary but fixed. First, we give an estimate for the probability that T is hosted by (ι, a, a^r) . By Lemma 2.6, for any $x, y \in (\Omega \setminus \Lambda)_{(|V(T)|)}$,

$$\text{Prob}(x^r = y \wedge r|_\Lambda = g) \geq (1 - \varepsilon') \frac{1}{n_{(|\Lambda| + |V(T)|)}} \text{ and } (1 + \varepsilon') \frac{1}{n_{(|\Lambda|)}} \geq \text{Prob}(r|_\Lambda = g).$$

Hence, for the conditional probability $\text{Prob}(x^r = y \mid r|_\Lambda = g)$, we have

$$(3.3.1) \quad \text{Prob}(x^r = y \mid r|_\Lambda = g) \geq \frac{1 - \varepsilon'}{1 + \varepsilon'} \frac{1}{(n - |\Lambda|)_{(|V(T)|)}} > \frac{1 - \varepsilon'}{1 + \varepsilon'} \frac{1}{n^{|V(T)|}}.$$

Since by definition T_α is hosted by (ι, a) , T is hosted by (ι, a, a^r) if and only if T_β is hosted by (ι, a^r) ; this is equivalent to

$$\begin{aligned} & (x\iota)^{a^r} = y\iota \text{ for all } (x, y) \in E(T_\beta) \\ \iff & (x\iota)^{r^{-1}ar} = y\iota \text{ for all } (x, y) \in E(T_\beta) \\ \iff & (x\iota)^{r^{-1}a} = (y\iota)^{r^{-1}} \text{ for all } (x, y) \in E(T_\beta) \\ \iff & \text{the function } V(T) \rightarrow \Omega, x \mapsto (x\iota)^{r^{-1}} \text{ is in } \mathcal{S}_\beta(T) \\ \iff & (V(T)\iota_z)^r = V(T)\iota \text{ for some } \iota_z \in \mathcal{S}_\beta(T). \end{aligned}$$

Lemma 2.3 implies that for different $\iota_{z_1}, \iota_{z_2} \in \mathcal{S}_\beta(T)$, the events $(V(T)\iota_{z_i})^r = V(T)\iota$ are disjoint. Therefore, also using (3.2.1) and (3.3.1),

$$\begin{aligned} \text{Prob}(T \text{ is hosted by } (\iota, a, a^r) \mid r|_\Lambda = g) &= \sum_{\iota_z \in \mathcal{S}_\beta(T)} \text{Prob}((V(T)\iota_z)^r = V(T)\iota \mid r|_\Lambda = g) \\ (3.3.2) \quad &> \frac{(1 - \varepsilon')^2}{1 + \varepsilon'} \frac{(1 - \delta)^{l_\beta(T)} \delta^{p_\beta(T)} n^{l_\beta(T) + p_\beta(T)}}{n^{|V(T)|}}. \end{aligned}$$

Next, by Lemma 2.3(ii) and (3.2.1), (3.3.2),

$$\begin{aligned} \mathbb{E}(|\text{fix}(w(a, a^r))| \mid r|_{\Lambda} = g) &\geq \sum_{j=1}^m \sum_{\iota \in \mathcal{S}_{\alpha}(T_j)} \text{Prob}(T_j \text{ is hosted by } (\iota, a, a^r) \mid r|_{\Lambda} = g) \frac{\text{fixed}(T_j)}{|\text{Aut}(T_j)|} \\ &> \sum_{j=1}^m \frac{(1 - \varepsilon')^3}{1 + \varepsilon'} \frac{(1 - \delta)^{l_{\alpha}(T_j) + l_{\beta}(T_j)} \delta^{p_{\alpha}(T_j) + p_{\beta}(T_j)} n^{l_{\alpha}(T_j) + l_{\beta}(T_j) + p_{\alpha}(T_j) + p_{\beta}(T_j)}}{n^{|V(T_j)|}} \frac{\text{fixed}(T_j)}{|\text{Aut}(T_j)|}. \end{aligned}$$

Finally, by Lemma 2.2, $l_{\alpha}(T_j) + l_{\beta}(T_j) + p_{\alpha}(T_j) + p_{\beta}(T_j) = |V(T_j)| + 1$, yielding

$$\mathbb{E}(|\text{fix}(w(a, a^r))| \mid r|_{\Lambda} = g) > (1 - \varepsilon) n \sum_{j=1}^m \frac{\delta_{T_j} \cdot \text{fixed}(T_j)}{|\text{Aut}(T_j)|}.$$

To finish the proof of the theorem, we simply take r that gives at least the expected number of fixed points. \square

4. PROOF OF THEOREM 1.3

In this section we apply Theorem 3.3 to prove Theorem 1.3. We start with two technical lemmas.

Lemma 4.1. *Let m be an integer and take $g, h \in \text{Sym}(m)$. In each of the following cases, $[h, (h^g)^{-1}][h, h^g]$ contains a 7-cycle:*

- (1) $m \geq 7$, $h = (1, 2, 3, \dots, m)$, g contains the cycle $(1, 3, m)$ and fixes the points $2, 4, 5, 6, m - 3, m - 2, m - 1$;
- (2) $m = 7$, $h = (1, 2, 3, 4, 5)(6)(7)$, $g = (1, 6)(3, 7)(2)(4)(5)$;
- (3) $m = 7$, $h = (1, 2, 3)(4, 5, 6)(7)$, $g = (1, 7, 2, 4)(3)(5)(6)$;
- (4) $m = 7$, $h = (1, 2)(3, 4)(5, 6)(7)$, $g = (1, 5, 7, 2, 3)(4)(6)$.

Proof. To show (1), we note simply that $(1, m, 5, 3, m - 1, 4, 2)$ is a 7-cycle of $[h, (h^g)^{-1}][h, h^g]$. Parts (2), (3) and (4) follow easily. \square

Note that in each case in Lemma 4.1 we prescribe the action of g on at most 10 points.

For $0 < \delta < 1$, define

$$\begin{aligned} f(\delta) := & (1 - \delta)^2 + (1 - \delta)^2 \delta + (1 - \delta)^3 \delta + 4(1 - \delta)^4 \delta^2 + 2(1 - \delta)^2 \delta^3 \\ & + 3(1 - \delta)^5 \delta^3 + 10(1 - \delta)^9 \delta^4 + 26(1 - \delta)^7 \delta^5 + 20(1 - \delta)^8 \delta^5 \\ (4.1.1) \quad & + 6(1 - \delta)^5 \delta^6 + 16(1 - \delta)^6 \delta^6 + 40(1 - \delta)^8 \delta^6 + 3(1 - \delta)^4 \delta^7 \\ & + 8(1 - \delta)^6 \delta^7 + 20(1 - \delta)^7 \delta^7 + 10(1 - \delta)^8 \delta^9 + 20(1 - \delta)^7 \delta^{10} \\ & + 10(1 - \delta)^6 \delta^{11} + 15(1 - \delta)^7 \delta^{11}. \end{aligned}$$

Lemma 4.2.

- (1) *The function $\delta \mapsto 1 - 0.999f(\delta)$ is monotone increasing on the interval $(0, 1)$.*
- (2) *The equation $0.999f(\delta) = 1 - \delta$ has a unique solution in $(0, 1)$. Up to six significant digits, the solution is $\delta = 0.632599$.*
- (3) *Starting with $\delta = 0.63$, nine iterations of the function $\delta \mapsto 1 - 0.999f(\delta)$ reach a value less than 0.326.*

Proof. All three results can be established using an algebra package such as Sage [S⁺09]. \square

We are now ready to prove Theorem 1.3.

Proof of Theorem 1.3. Let a be an element of S with $|\text{supp}(a)| < Cn = 0.63n$. We shall apply the results of Section 3 for the word $w = w_0^{60}$, where $w_0 = w_0(\alpha, \beta) = [\alpha, \beta^{-1}][\alpha, \beta]$.

The proof splits into a number of cases, according to the order of a . Let $|a| = 2^{e_1} 3^{e_2} e_3$, with e_3 coprime to 6. Note that $2^{e_1} \leq n$ and $3^{e_2} \leq n$. Suppose first that $e_3 > 1$. Then $a^{2^{e_1} 3^{e_2}}$ is non-trivial, and we may replace a with $a^{2^{e_1} 3^{e_2}}$ and assume that the order of a is coprime to both 2 and 3 (note that $a^{2^{e_1} 3^{e_2}} \in (S \cup S^{-1} \cup \{1\})^{n^2}$).

In Tables 1, 2, 3, 4, and 5, we present a family \mathcal{T} of $\alpha\beta$ -trees admitting w_0^{60} . These $\alpha\beta$ -trees were obtained with the help of a computer; note that our conventions for representing them are outlined in Subsection 2.3. Note too that if an $\alpha\beta$ -graph admits w_0^e then it also admits w_0^{ek} for any positive integer k . The $\alpha\beta$ -trees

in these tables are pairwise non-isomorphic, do not have non-identity automorphisms, contain at most 16 vertices, and have all α -paths and β -paths of length at most 4. For this family,

$$\sum_{T \in \mathcal{T}} \frac{\delta_T \cdot \text{fixed}(T)}{|\text{Aut}(T)|} = f(\delta),$$

where $f(\delta)$ is as in (4.1.1). If a has a cycle of length $m \geq 7$ then relabel the letters of Ω so that this cycle is equal to $(1, \dots, m)$, define $\Lambda = \{1, 2, 3, 4, 5, 6, m-3, m-2, m-1, m\}$ and let $g = (1, 3, m)(2)(4)(5)(6)(m-3)(m-2)(m-1) \in \text{Sym}(\Lambda)$. If this is not the case, then we must have $e_3 = 5$; in this case relabel the letters of Ω so that a contains the cycle $(1, 2, 3, 4, 5)$ and a fixes both 6 and 7. Then define $\Lambda = \{1, \dots, 7\}$ and let $g = (1, 6)(3, 7)(2)(4)(5)$.

By Lemma 2.1, we may suppose that $|S| \leq n$. We apply Theorem 3.3 with $\delta_0 = 0.3$, $w = w_0^{60}$, \mathcal{T} , $N = 4$, $\lambda = 10$, g , $\kappa = 16$, and $\varepsilon = 0.001$. We obtain $r \in (S \cup S^{-1} \cup \{1\})^{n^{54}}$ such that $|\text{fix}(w(a, a^r))| \geq 0.999f(\delta)$. Note that $w(a, a^r) \in (S \cup S^{-1} \cup \{1\})^{n^{54}+n^2}$ and, of course, $n^{54} + n^2 = O(n^{54})$. By Lemma 4.1, (1) and (2), $w(a, a^r)$ is non-trivial because it contains a 7-cycle. Replacing a by $w(a, a^r)$ and using the same procedure as above, Lemma 4.2, part (3), implies that in at most nine iterations we obtain a permutation a' with support size less than $0.326n$. Each iteration increases the word length by a factor $O(n^2)$, as we may have to raise the input permutation to a suitable power to eliminate 2 and 3 from the cycle lengths, while conjugating by (a new) r and substituting into the word w contributes only constant multipliers to the word length. Hence a' is a word in S of length $O(n^{70})$ and, by Theorem 1.2, $\text{diam}(\Gamma(G, S)) = O(n^{78})$.

Suppose next that $e_3 = 1$, that is, a has order $2^{e_1}3^{e_2}$. If $e_1 > 0$, let $k = 2^{e_1-1}3^{e_2}$, otherwise let $k = 3^{e_2-1}$. Then $k < n^2$, a^k has order 2 or 3, and $a^k \in (S \cup S^{-1} \cup \{1\})^{n^2}$.

In Tables 6 and 7, we present two families of $\alpha\beta$ -trees admitting $w = w_0^{60}$, represented as before. The $\alpha\beta$ -trees in Table 6 (resp. Table 7) are pairwise non-isomorphic, contain at most 10 vertices, have all α -cycles and β -cycles of length 2 (resp. 3), and have all α -paths and β -paths of length at most 1 (resp. 2).

In each cell of Tables 6 and 7, we have written $\text{Aut} = k$ to mean that the automorphism group of the corresponding $\alpha\beta$ -tree T has size k . We define \mathcal{T} to be the set of $\alpha\beta$ -trees in Table 6 (resp. Table 7) when a has order 2 (resp. has order 3). For these families,

$$\sum_{T \in \mathcal{T}} \frac{\delta_T \cdot \text{fixed}(T)}{|\text{Aut}(T)|} = h(\delta)$$

where

$$h(\delta) = \begin{cases} (1 - \delta)^2(1 + 2\delta + 3\delta^2 + 4\delta^3 + 5\delta^4 + 6\delta^5) & \text{if } a \text{ has order 2;} \\ \begin{aligned} &(1 - \delta)^2 + \delta(1 - \delta)^2 + \delta(1 - \delta)^3 + 2\delta^3(1 - \delta)^2 \\ &+ 4\delta^2(1 - \delta)^4 + 3\delta^3(1 - \delta)^5 + 12\delta^7(1 - \delta)^4 \\ &+ 6\delta^6(1 - \delta)^4 + \delta^4(1 - \delta)^6 \end{aligned} & \text{if } a \text{ has order 3.} \end{cases}$$

Define $\Lambda = \{1, \dots, 7\}$. If a has order 2 label the elements of Ω so that $a|_\Lambda = (1, 2)(3, 4)(5, 6)(7)$; then define $g = (1, 5, 7, 2, 3)(4)(6)$. If a has order 3 label the elements of Ω so that $a|_\Lambda = (1, 2, 3)(4, 5, 6)(7)$; then define $g = (1, 7, 2, 4)(3)(5)(6)$.

By Lemma 2.1, we may suppose that $|S| \leq n$. We define N to equal 1 (resp. 2) when a has order 2 (resp. has order 3). We apply Theorem 3.3 with $\delta_0 = 0.3$, $w = w_0^{60}$, \mathcal{T} , N , $\lambda = 7$, g , $\kappa = 10$, and $\varepsilon = 0.001$. We obtain $r \in (S \cup S^{-1} \cup \{1\})^{n^{36}}$ such that $|\text{fix}(w(a, a^r))| \geq 0.999h(\delta)$.

Using Sage [S⁺09] it is easy to check that the function $\delta \mapsto 1 - 0.999h(\delta)$ is monotone increasing on the interval $(0, 1)$. Furthermore, for $\delta \leq 0.63$, we have $0.999h(\delta) > 0.374$ and so $|\text{supp}(w(a, a^r))| < 0.626$. Note that $w(a, a^r) \in (S \cup S^{-1})^{O(n^{36})}$ and, by Lemma 4.1, (3) and (4), $w(a, a^r)$ contains a 7-cycle.

We now run the first part of the argument using the element $w(a, a^r)$ instead of a as our initial element of small support. After one iteration we obtain an element $a' \in (S \cup S^{-1} \cup \{1\})^{n^{2O(n^{36})+n^{54}}}$ with support of size less than $1 - 0.999f(\delta)$. Iterating as before, we obtain that $\text{diam}(\Gamma(G, S)) = O(n^{78})$. \square

5. IMPROVING THEOREM 1.3

It should be clear to the reader that Theorem 1.3 is not optimal. In particular we prove Theorem 1.3 with respect to a particular word, $w = w_0^{60}$ where $w_0 = [\alpha, \beta^{-1}][\alpha, \beta]$; it is this word that yields the value $C = 0.63$. How might one go about improving this value?

The most obvious way of improving Theorem 1.3 is via an appeal to higher powers. Consider $w_k = w_0^k$ where k is any multiple of 60. All of the trees in Tables 1, 2, 3, 4 and 5 admit w_k and, for suitable choices of k , there will be yet more trees to consider. This will inevitably result in an increase for the value of C .

There is a limit to the improvement that such a strategy might yield and we briefly explain why this is the case. Fix a word v_0 , let k be some positive integer, and define the word $v_k = v_0^k$. Let \mathcal{T} be a set of $\alpha\beta$ -trees admitting v_k and consider the sum

$$(5.0.1) \quad g_k(\delta) = \sum_{T \in \mathcal{T}} \frac{\delta_T \cdot \text{fixed}(T)}{|\text{Aut}(T)|}.$$

The main result of Section 3, Theorem 3.3, gives a lower bound for $|\text{fix}(w(a, a^r))|$ in terms of $g_k(\delta)$, ε and n ; here a is an element of support equal to δn and r is some short word. In particular, if the value of $g_k(\delta)$ exceeds the value of $1 - \delta$ (by a sufficient margin in terms of ε) then we obtain an element of smaller support than a .

The advantage of considering higher powers of the word v_0 is exhibited in Theorem 3.3 by noting that if, say, k doubles, then new $\alpha\beta$ -trees may be added to \mathcal{T} , thereby increasing the value of $g_k(\delta)$ for $\delta \in (0, 1)$.

Using methods different to those in this paper, the authors have developed a method to show that, in the “generic” case (see Section 3 for an explanation of what we mean by this), there is a number $\delta_0 \in (0, 1)$ such that

$$\limsup_{k \rightarrow \infty} g_k(\delta) < 1 - \delta$$

whenever $\delta > \delta_0$. It is important to note that the number δ_0 depends only on the word v_0 . Furthermore the number δ_0 corresponds to the unique all-positive solution to a certain system of polynomial equations with rational coefficients and this system depends only on the word v_0 .

In the case $v_0 = w_0 = [\alpha, \beta^{-1}][\alpha, \beta]$, our method shows that $\delta_0 \approx 0.64242$. One can see, then, that the value for C given in Theorem 1.3 is close to optimal when it comes to powers of the word w_0 .

In another direction one might hope to improve Theorem 1.3 using an entirely different choice of word w_0 . With this in mind we undertook an exhaustive search of words of length at most 20 which were *balanced* (i.e. $\alpha, \alpha^{-1}, \beta$ and β^{-1} all occur the same number of times) and in which $\alpha^{\pm 1}$ and $\beta^{\pm 1}$ occur alternately. For each such word w_1 , we performed the following computer experiment. For a variety of values n in the range $10^4 \leq n \leq 10^5$ and δ in the range $0.55 \leq \delta \leq 0.70$, we took a random permutation $a \in S_n$ with $|\text{supp}(a)| = \delta n$, we constructed b as a random conjugate of a , and counted how many points occur in cycles of length at most 6 in the permutation $w_1(a, b)$. The highest counts occurred for $w_0 = [\alpha, \beta^{-1}][\alpha, \beta]$ and for related words (like the cyclic permutations of w_0 or w_0^2) and this is the reason for our use of the word w_0 in the preceding proof.

We also carried out a non-exhaustive investigation into words that were either non-balanced or non-alternating. In every case, for a word w_1 of this kind, and for a and b as above, the computer tests suggested that permutations of the form $w_1(a, b)$ tended to have a smaller number of points in short cycles than $w_0(a, b)$.

REFERENCES

- [1] L. Babai, On the diameter of Eulerian orientations of graphs, *Proc. of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ACM, New York, 2006, 822–831.
- [2] L. Babai, R. Beals, Á. Seress, On the diameter of the symmetric group: polynomial bounds, *Proc. of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ACM, New York, 2004, 1108–1112.
- [3] L. Babai, Á. Seress. On the diameter of Cayley graphs of the symmetric group, *J. Combin. Theory Ser. A*, **49** (1988), 175–179.
- [4] L. Babai, Á. Seress, On the diameter of permutation groups, *European J. Combin.* **13** (1992), 231–243.
- [5] A. Bochert, Über die Classe der transitiven Substitutionengruppen (German), *Math. Ann.* **49** (1897), 133144.
- [6] E. Breuillard, B. Green, T. Tao, Linear approximate groups, *Geometric and Functional Analysis*, to appear, Preprint available on the Math arXiv: <http://arxiv.org/abs/1001.4570>.
- [7] P. Erdős, A. Rényi, Probabilistic methods in group theory, *J. Analyse Math.* **14** (1965), 127–138.
- [8] N. Gill, H. A. Helfgott, Growth of small generating subsets in $SL_n(\mathbb{Z}/p\mathbb{Z})$, *Int. Math. Res. Not.* **18** (2011), 4226–4251.
- [9] ———, Growth in solvable subgroups of $GL_r(\mathbb{Z}/p\mathbb{Z})$, 2010. Preprint available on the Math arXiv: <http://arxiv.org/abs/1008.5264>.
- [10] N. Gill, H. A. Helfgott, Á. Seress, P. Spiga. Preprint in preparation.
- [11] H. A. Helfgott, Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$, *Ann. of Math. (2)* **167** (2008), 601–623.
- [12] ———, Growth and generation in $SL_3(\mathbb{Z}/p\mathbb{Z})$, *J. Eur. Math. Soc.* **13** (2011), no. 3, 761851.

- [13] H. Helfgott, Á. Seress, On the diameter of permutation groups, 2011, Preprint available on the math arXiv: <http://arxiv.org/abs/1109.3550>.
- [14] H. J. Landau, A. M. Odlyzko, Bounds for eigenvalues of certain stochastic matrices, *Linear algebra and its Applications* **38** (1981), 5–15.
- [15] L. Lovász, Random walks on graphs: a survey, *Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993)*, volume 2 of *Bolyai Soc. Math. Stud.*, János Bolyai Math. Soc., Budapest, 1996, 353–397.
- [16] W. A. Manning, The degree and class of multiply transitive groups, *Trans. Amer. Math. Soc.* **18** (1917), no. 4, 463–479.
- [17] ———, The degree and class of multiply transitive groups. II, *Trans. Amer. Math. Soc.* **31** (1929), no. 4, 643–653.
- [18] ———, The degree and class of multiply transitive groups. III, *Trans. Amer. Math. Soc.* **35** (1933), no. 3, 585–599.
- [19] C. E. Praeger, L. Pyber, P. Spiga, E. Szabo, The Weiss conjecture for locally primitive graphs with automorphism groups admitting composition factors of bounded rank, To appear in *Proc. Amer. Math. Soc.*
- [20] L. Pyber, E. Szabó, Growth in finite simple groups of Lie type, 2010, Preprint available on the Math arXiv: <http://arxiv.org/abs/1001.4556>.
- [21] E. Z. Ruzsa, Sums of finite sets, In *Number theory*, Springer, New York, 1996, 281–293.
- [22] Á. Seress, *Permutation group algorithms*, Cambridge Tracts in Mathematics, **152**, Cambridge University Press, 2003.
- [23] P. Spiga, Two local conditions on the vertex stabiliser of arc-transitive graphs and their effect on the Sylow subgroups *J. Group Theory* **15**, no. 1 (2012), 23–35.
- [S⁺09] W. A. Stein et al., *Sage Mathematics Software (Version 4.8)*, The Sage Development Team, 2012, <http://www.sagemath.org>
- [24] J. Whiston, Maximal independent generating sets of the symmetric group, *J. Algebra* **232** (2000), no. 1, 255–268.
- [25] H. Wielandt, Abschätzungen für den Grad einer Permutationsgruppe von vorgeschriebenem Transitivitätsgrad, Dissertation, (1934) Berlin.




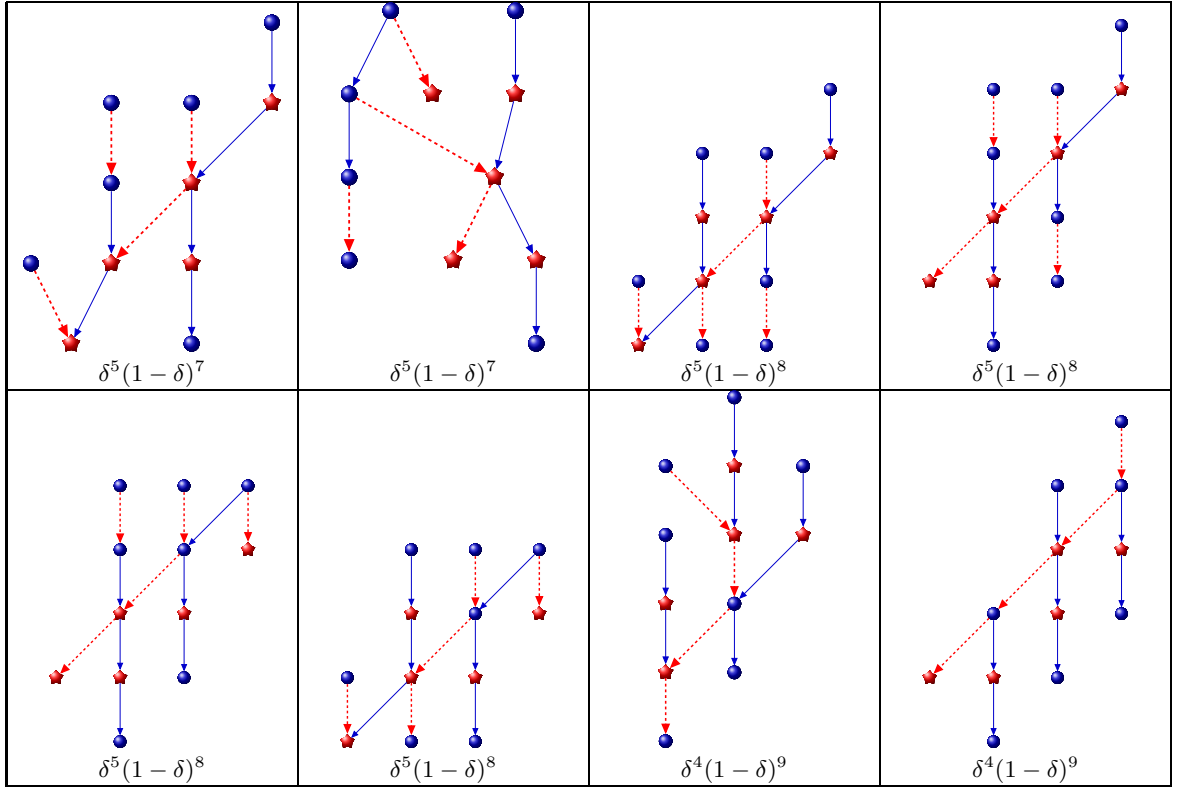
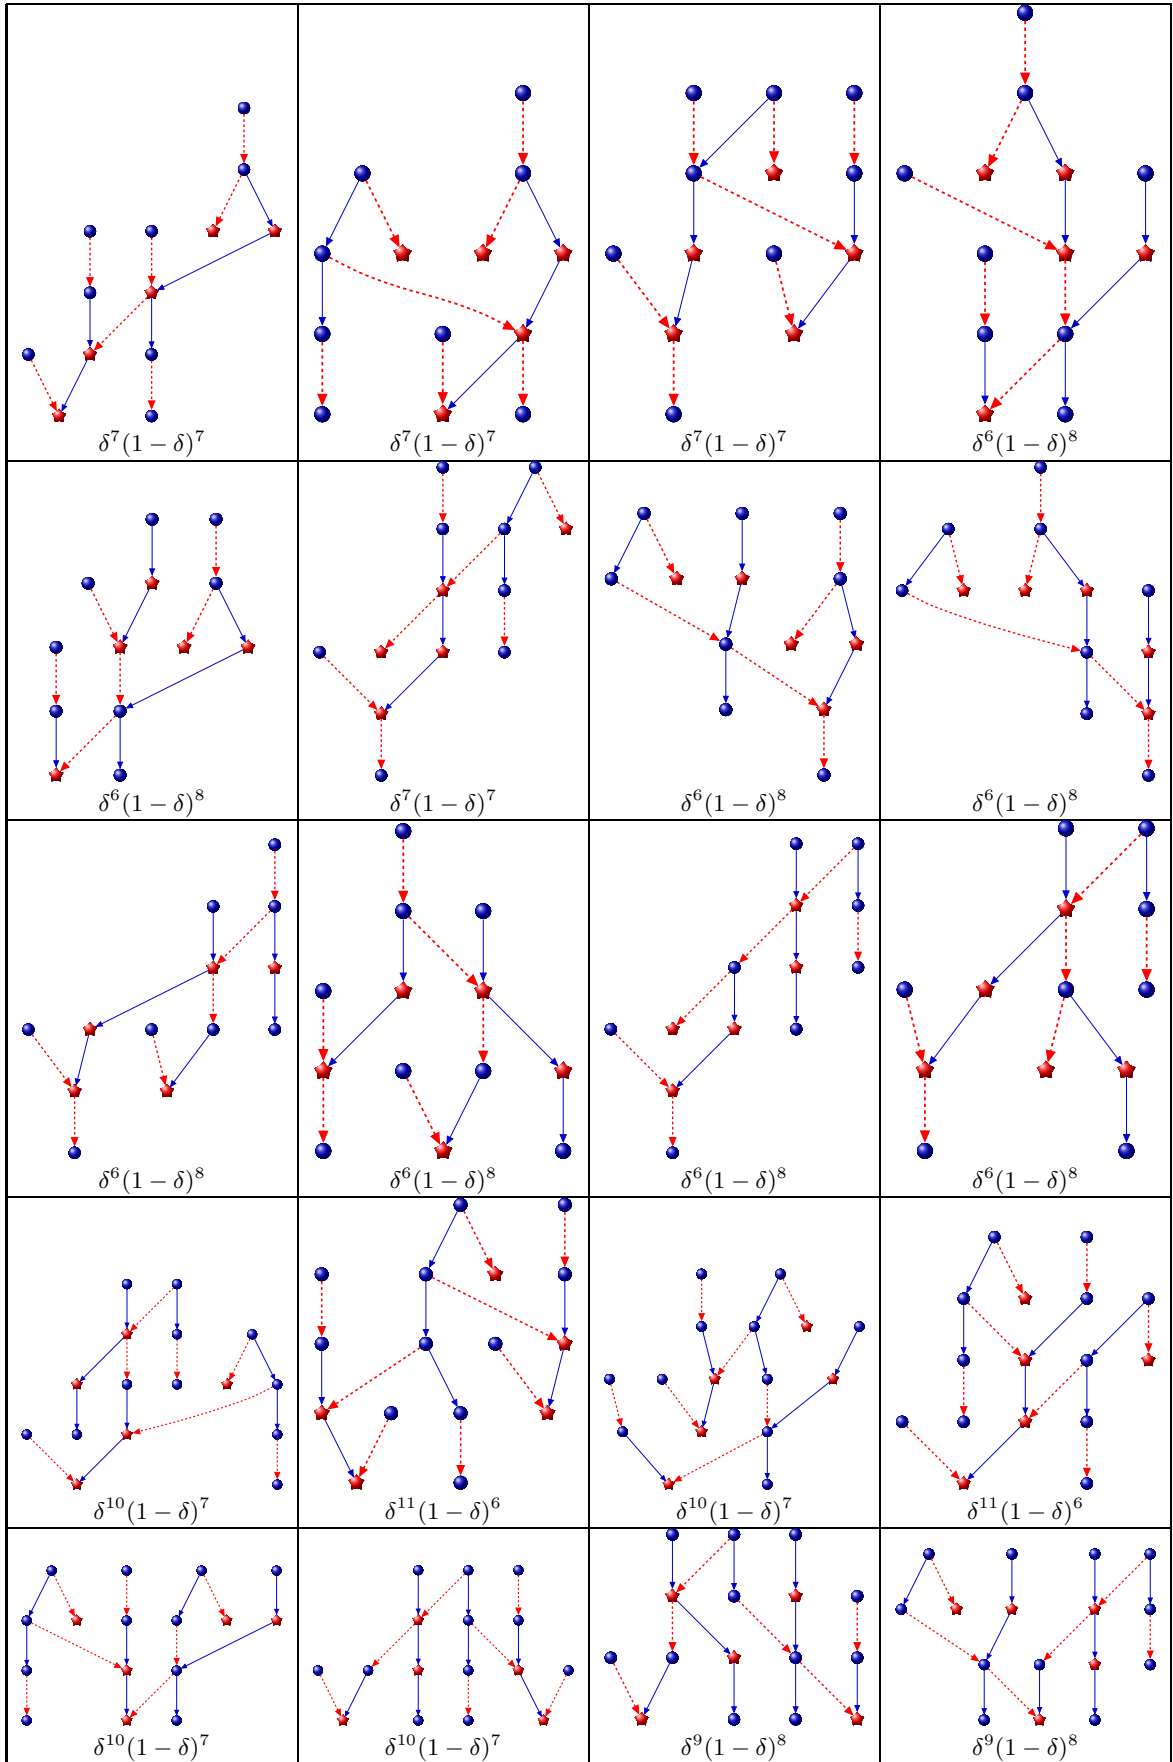


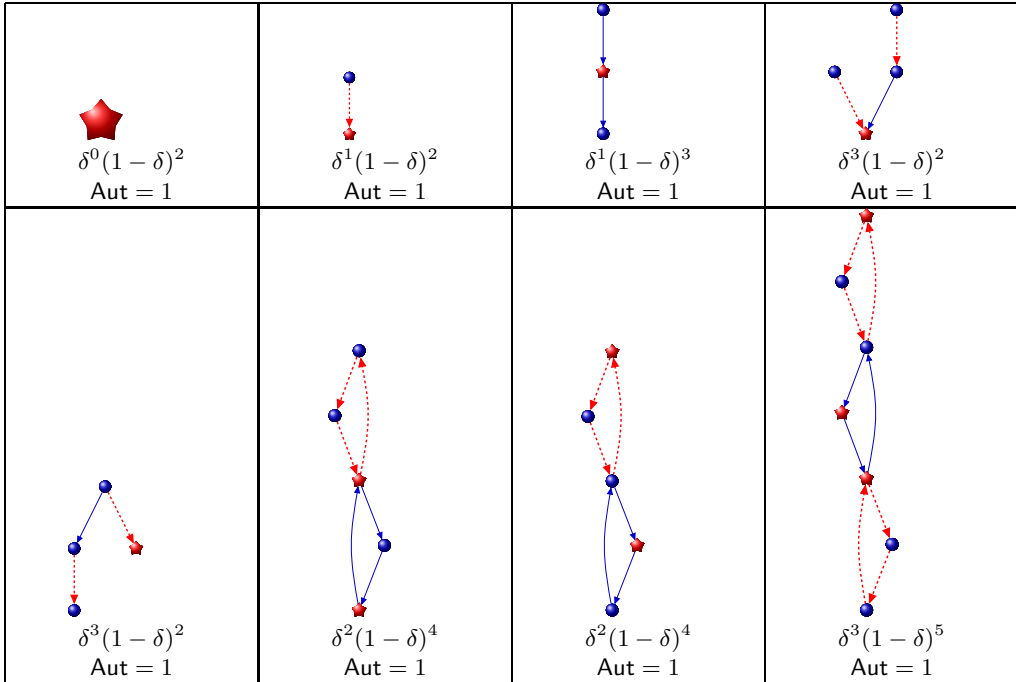

TABLE 2. $\alpha\beta$ -trees admitting w^2 but not w TABLE 3. $\alpha\beta$ -trees admitting w^3 but not w 

TABLE 4. $\alpha\beta$ -trees admitting w^4 but not w^2

 Table 5: $\alpha\beta$ -trees admitting w^5 but not w





 Table 6: $\alpha\beta$ -trees for an involution

 Table 7: $\alpha\beta$ -graphs for an element of order 3




JOHN BAMBERG
 SCHOOL OF MATHEMATICS AND STATISTICS
 UNIVERSITY OF WESTERN AUSTRALIA
 35 STIRLING HIGHWAY, CRAWLEY, WA 6009, AUSTRALIA
E-mail address: john.bamberg@uwa.edu.au

NICK GILL
 DEPARTMENT OF MATHEMATICS
 THE OPEN UNIVERSITY
 WALTON HALL, MILTON KEYNES, MK7 6AA, UNITED KINGDOM
E-mail address: n.gill@open.ac.uk

THOMAS P. HAYES
 DEPARTMENT OF COMPUTER SCIENCE
 MAIL STOP: MSC01 1130
 1 UNIVERSITY OF NEW MEXICO
 ALBUQUERQUE, NM 87131-0001
E-mail address: hayes@cs.unm.edu

HARALD A. HELFGOTT
 DÉPARTEMENT DE MATHÉMATIQUES ET APPLICATIONS
 ÉCOLE NORMALE SUPÉRIEURE
 45 RUE D'ULM
 F-75230 PARIS, FRANCE
E-mail address: helfgott@dma.ens.fr

ÁKOS SERESS
 SCHOOL OF MATHEMATICS AND STATISTICS
 UNIVERSITY OF WESTERN AUSTRALIA
 35 STIRLING HIGHWAY, CRAWLEY, WA 6009, AUSTRALIA
 AND
 DEPARTMENT OF MATHEMATICS
 THE OHIO STATE UNIVERSITY
 231 W. 18TH AVENUE, COLUMBUS, OH 43210, USA
E-mail address: akos@math.ohio-state.edu

PABLO SPIGA
DIPARTIMENTO DI MATEMATICA E APPLICAZIONI
UNIVERSITY OF MILANO-BICOCCA
VIA COZZI 53, 20125 MILANO, ITALY
E-mail address: `pablo.spiga@unimib.it`